

Efficient method for measuring the parameters encoded in a gravitational-wave signal

Carl-Johan Haster

School of Physics and Astronomy, University of Birmingham, Edgbaston,
Birmingham B15 2TT, United Kingdom

Ilya Mandel

School of Physics and Astronomy, University of Birmingham, Edgbaston,
Birmingham B15 2TT, United Kingdom
Monash Center for Astrophysics, Monash University, Clayton, VIC 3800, Australia

Will M. Farr

School of Physics and Astronomy, University of Birmingham, Edgbaston,
Birmingham B15 2TT, United Kingdom

E-mail: cjhaster@star.sr.bham.ac.uk; imandel@star.sr.bham.ac.uk;
w.farr@bham.ac.uk

Abstract. Once upon a time, predictions for the accuracy of inference on gravitational-wave signals relied on computationally inexpensive but often inaccurate techniques. Recently, the approach has shifted to actual inference on noisy signals with complex stochastic Bayesian methods, at the expense of significant computational cost. Here, we argue that it is often possible to have the best of both worlds: a Bayesian approach that incorporates prior information and correctly marginalizes over uninteresting parameters, providing accurate posterior probability distribution functions, but carried out on a simple grid at a low computational cost, comparable to the inexpensive predictive techniques.

Submitted to: *Class. Quantum Grav.*

PACS numbers: 02.50.Tt, 04.30.-w, 95.85.Sz

1. Introduction

Advanced LIGO gravitational-wave (GW) detectors are expected to start science observations in 2015, with Advanced Virgo following shortly after [1, 2, 3, 4, 5]. A prime source predicted to be observable by these detectors are compact binary coalescences (CBCs) during which neutron star (NS) or black hole (BH) components of the binary are driven by GW emission through an inspiral into a merger. Binary neutron stars (BNS) have already been observed as pulsars in Galactic binary systems evolving due to emission of GWs [6, 7]; extrapolations from these observations yield estimated coalescence rates of BNS systems as detectable by the Advanced LIGO–Virgo network ([8] and references therein).

1.1. Binary coalescence model

The information available in an observed circular CBC event is encoded through the waveform \vec{h} , fully described by a 15-dimensional parameter vector $\vec{\theta}$ which in turn can further be divided into two parameter sets. The amplitude and phase evolution of the waveform are governed by the intrinsic parameters which include the masses and spins of the binary components. The extrinsic parameters describe the projection of the gravitational wave emitted by the binary onto the detectors and include the object’s position and orientation on the sky, its luminosity distance d_L as well as the time and phase of the waveform at coalescence (t_c and ϕ_c). The level of correlation between the two parameter sets is only marginal, apart for t_c and ϕ_c which are strongly correlated with intrinsic parameters. Therefore we will simplify the analysis by fixing the remaining extrinsic parameters at their true values without significant impact on the recovery of the intrinsic parameters; we return to this point in section 4.

In this study we consider as an example a BNS with component masses $m_1 = 1.45M_\odot$ and $m_2 = 1.35M_\odot$. While there is a strong degeneracy between the mass and spin parameters, which affects their measurement uncertainty [9, 10, 11, 12, 13], the BNS components are in this study taken as non-spinning for demonstration purposes, as our focus here is on parameter estimation techniques rather than astrophysical predictions. Therefore, the original 15-dimensional parameter space is reduced to only four parameters, the two masses as well as t_c and ϕ_c , without loss of generality in the method applied. The waveform approximant used is TaylorF2 [14] which for a face-on binary directly overhead the detector describes the frequency-domain waveform as

$$\tilde{h}(f, \vec{\theta}) = \frac{1}{d_L} \sqrt{\frac{5}{24}} \pi^{-2/3} \mathcal{M}_c^{5/6} f^{-7/6} e^{i\Psi(f, \vec{\theta})}, \quad (1)$$

assuming no higher-order corrections to the waveform amplitude. Here, $\mathcal{M}_c \equiv (m_1 m_2)^{3/5} (m_1 + m_2)^{-1/5}$ is the chirp mass and the phase $\Psi(f, \vec{\theta})$ is given as

$$\Psi(f, \vec{\theta}) = 2\pi f t_c - \phi_c - \frac{\pi}{4} + \frac{3}{128\eta\nu^5} \sum_{k=0}^7 (\alpha_k + \beta_k \log(\nu)) \nu^k \quad (2)$$

where $\nu \equiv (\pi(m_1 + m_2)f)^{1/3}$ and the coefficients α_k, β_k are functions of \mathcal{M}_c , the spins (here assumed to be zero) and the mass ratio $\eta \equiv (m_1 m_2)(m_1 + m_2)^{-2}$ for each post-Newtonian order $k/2$.

The likelihood of observing data $d \equiv h_T + n$, containing the true waveform h_T and stationary, Gaussian instrumental noise n , given a waveform model $h(\vec{\theta})$ described by parameters $\vec{\theta}$ is

$$L(\vec{\theta}) \propto \exp \left(-\frac{1}{2} \langle n + h_T - h(\vec{\theta}) | n + h_T - h(\vec{\theta}) \rangle \right). \quad (3)$$

The term in the angle brackets is the noise-weighted power in the residuals after subtracting the assumed model from the data. The angle brackets denote an inner product between the Fourier transformed time series $\tilde{a}(f)$ and the complex conjugate of $\tilde{b}(f)$ weighted by the noise power spectral density $S_n(f)$:

$$\langle a | b \rangle = 4 \operatorname{Re} \int_0^\infty \frac{\tilde{a}(f) \tilde{b}^*(f)}{S_n(f)} df. \quad (4)$$

The likelihood (3), together with the specification for the waveform model and the noise power spectral density $S_n(f)$, which is chosen to match the Advanced LIGO zero-detuning high-power design sensitivity [15], constitutes the model H .

For this study we will consider a BNS observed directly overhead in one detector at a signal to noise ratio $\rho = 12$, defined as in [16]:

$$\rho \equiv \max_{\vec{\theta}} \frac{\langle d | h(\vec{\theta}) \rangle}{\sqrt{\langle h(\vec{\theta}) | h(\vec{\theta}) \rangle}}. \quad (5)$$

The injected noise realisation in our example is chosen to be $n \equiv \vec{0}$, which fixes the maximum likelihood waveform to be h_T .

1.2. Bayesian inference

Given the model H as described in subsection 1.1, (3) can be viewed as a probability density function (PDF) $p(d|\vec{\theta}, H)$ for observing data $d \equiv h_T + n$ given parameters $\vec{\theta}$. Together with a PDF on the parameters detailing the prior knowledge and expectation about $\vec{\theta}$ under H , $p(\vec{\theta}|H)$, a posterior PDF can be obtained through Bayes' theorem as

$$p(\vec{\theta}|d, H) = \frac{p(\vec{\theta}|H)p(d|\vec{\theta}, H)}{p(d|H)}. \quad (6)$$

The term in the denominator is called the evidence and can be used for comparing the effectiveness of different models' ability to describe the data; for this study, it is a normalizing constant for the posterior PDF. The task of parameter estimation can therefore be defined as finding accurate and robust methods for evaluating the posterior probability given by (6) in order to obtain the accuracy with which different parameters can be recovered. This accuracy is usually quantified in terms of credible regions CR_p

containing a fraction p of the total posterior probability. These credible regions are defined as

$$\text{CR}_p \equiv \min A \tag{7a}$$

such that the parameter-space area A satisfies

$$p = \int_A p(\vec{\theta}|d, H) d\vec{\theta}. \tag{7b}$$

The posterior PDF will have the same dimensionality as $\vec{\theta}$ and a shape strongly dependent on the complexity of H . For the PDFs encountered in CBC data analysis, very intricate structures with high levels of multimodality and parameter degeneracies are common, and their prevalence has spurred the development and use of advanced analysis tools for sampling the posterior PDF.

1.3. Stochastic sampling

As part of the publicly available LSC Algorithm Library (LAL) [17], the Bayesian framework in **LALInference** implements several methods to stochastically traverse the parameter space [18, 19], creating a set of individual samples $\vec{\theta}_i$ distributed according to the posterior PDF in (6). For this study, we used a parallel-tempered Markov chain Monte Carlo (MCMC) algorithm [20, 21], implemented in **LALInference** as **LALInferenceMCMC**. The number of MCMC samples required to fully describe the posterior PDF scales only weakly with the number of dimensions, allowing for efficient exploration of high-dimensional parameter spaces. However, only a small fraction of the collected samples will be statistically independent, due to imperfect jump proposals in the complex multimodal parameter spaces. In order to accurately represent the posterior PDF for the 4-dimensional parameter vector $\vec{\theta}$ from subsection 1.1, $\mathcal{O}(10^3)$ independent samples are necessary [22]. This requires a total of $\mathcal{O}(10^6)$ **LALInferenceMCMC** posterior samples, and therefore $\mathcal{O}(10^7)$ individual likelihood calculations spread across the 9 differently tempered parallel chains used in this example.

1.4. Paper organisation

In section 2 we present the “cumulative marginalized posterior” method for accurately and efficiently evaluating credible region contours on a grid. We show that credible regions can be computed with a suitable choice of a low-density grid of samples, described in subsection 2.2. This makes it possible to produce credible regions with the same accuracy as given by the stochastic samplers, but at a greatly reduced computational cost, as shown in subsection 2.3. In section 3 we discuss alternative grid-based methods for estimating credible regions, including the iso-match contour technique advocated by Baird *et al.* [23], and show their relative performance against the cumulative marginalized posterior method (subsection 3.3). We conclude in section 4 with a discussion of these results and a proposal for implementing the cumulative

marginalized posterior method in a low-latency parameter estimation framework for GW signals from CBCs.

2. Discretizing the credible regions

As an alternative to the stochastic sampling methods we propose a method to evaluate the integral in (7b) using samples at pre-determined coordinates in parameter space. The simplest implementation places samples into centres of pixels distributed in a uniform rectangular grid and approximates the posterior as constant within a pixel. The choice of the sampling grid is critical to the efficiency of this method, and will be discussed in subsection 2.2.

The required number of likelihood calculations, and therefore the required number of pixels, scales exponentially with the number of dimensions in the grid, so in order to minimize the computational requirements a modified version of (3) was implemented. We use a likelihood function marginalized over t_c and ϕ_c , thus removing the need for sampling those dimensions without affecting the PDF in the mass parameters of interest. This marginalized likelihood function is presented in [24], which also provides constraints on the resolution at which the waveforms must be sampled in both the frequency and time domains, in order to guarantee no loss of information when computing the inner products.

2.1. Cumulative posterior on a grid

Taking advantage of the natural $[\mathcal{M}_c, \eta]$ parameterization of the waveforms, we constructed a uniformly spaced rectangular grid across these parameters. For each pixel in the grid, the likelihood corresponding to the value of $\vec{\theta}$ at its midpoint was calculated. We assumed a prior PDF on the mass distribution of BNS systems to be uniform in $[m_1, m_2]$ rather than $[\mathcal{M}_c, \eta]$ (transformed appropriately when working in the $[\mathcal{M}_c, \eta]$ space). The product of the likelihood and prior yield the numerator of (6). The posterior PDF is obtained by normalizing this quantity by the evidence, the denominator in (6), which is approximated as the sum of the likelihood-prior products over all pixels in the grid. The posterior PDF is assumed to be valid not just locally at $\vec{\theta}$ but instead across a whole pixel. Credible regions can then be defined by the set of pixels containing a fraction p of the posterior PDF, accumulated when traversing the pixels in order of decreasing posterior values.

2.2. Grid placement

To minimize computational cost and to ensure an accurate representation of the credible regions defined by the integral in (7b), the grid samples must be placed as sparsely as possible, subject to two constraints: (i) a sufficient fraction of the parameter space with significant posterior support is covered to enable accurate normalization of the posterior; and (ii) the error introduced by approximating the prior-likelihood integral

over any pixel as the product of the prior and likelihood at the centre and the pixel area is within desired bounds.

The required extent of the grid can be quantified in terms of the Mahalanobis distance $r(\vec{\theta})$ defined as

$$r = \sqrt{(\vec{\theta} - \vec{\mu})\Sigma^{-1}(\vec{\theta} - \vec{\mu})^T} \quad (8)$$

for a set of pixel coordinates $\vec{\theta}$ spanning a multivariate N -dimensional PDF $f(\vec{\theta}) \equiv p(\vec{\theta}|d, H)$ with mean $\vec{\mu}$ and covariance Σ [25]. When $f(\vec{\theta})$ is a bivariate Gaussian, the associated cumulative density function is given by $\Phi(r) = 1 - e^{-r^2/2}$. Hence, for a maximum error $\epsilon = 1 - \Phi(r)$ in the evidence contained within the grid, the grid must minimally contain the pixels bounded by a distance

$$r_b = \sqrt{-2 \ln \epsilon} \quad (9)$$

away from the maximum (mean) of the PDF. The main purpose of this analysis is to construct credible regions whose p -value is known to an accuracy no worse than that of a stochastic sampler, which is of order 1% for the $\mathcal{O}(10^3)$ samples we typically have (see subsection 2.3). We therefore set $\epsilon = 0.005$, and will correspondingly cover the region $r_b \leq 3.25$ with a grid.

The minimum density of pixels within this bound is set by requiring that the approximate PDF, computed by discretely evaluating the posterior on a grid, is a sufficiently good approximation to the integral (7b). Here we use a very simple approximation, namely, we evaluate the integral as a Riemann sum over equal-sized pixels, setting the contribution of each pixel to the integral equal to the product of the pixel area and the value of the PDF at the centre of the pixel. In this case, a sufficient – but not necessary – condition for the total error on the PDF integral to be bounded by ϵ is for the fractional error in each pixel to be smaller than ϵ . For an N -dimensional PDF $f(\vec{\theta})$, this fractional difference across a pixel centred at $\vec{\theta}_0$ is

$$\left| \frac{f(\vec{\theta}_0)\Delta^N - \int_{\vec{\theta}_0 - \vec{\Delta}/2}^{\vec{\theta}_0 + \vec{\Delta}/2} f(\vec{\theta}) d\vec{\theta}}{f(\vec{\theta}_0)\Delta^N} \right| \leq \epsilon, \quad (10)$$

where Δ is the pixel width. For a one-dimensional Gaussian distribution $f(\theta)$ with zero mean and variance σ^2 , the integral in (10) can be represented as a Taylor series:

$$\begin{aligned} & \int_{\theta_0 - \Delta/2}^{\theta_0 + \Delta/2} f(\theta) d\theta \\ &= \int_{\theta_0 - \Delta/2}^{\theta_0 + \Delta/2} f(\theta_0) \left(1 - \frac{\theta_0}{\sigma^2}(\theta - \theta_0) - \frac{1}{2\sigma^2}(\theta - \theta_0)^2 + \frac{\theta_0^2}{2\sigma^4}(\theta - \theta_0)^2 + \mathcal{O}((\theta - \theta_0)^3) \right) d\theta \quad (11) \\ &\approx f(\theta_0) \left(\Delta - \frac{\Delta^3}{24\sigma^2} + \frac{\theta_0^2 \Delta^3}{24\sigma^4} \right) \end{aligned}$$

where the first non-zero correction term enters at the second order of the Taylor series since the approximated PDF is evaluated at the centre of the pixel. The integral will be dominated by the last term in (11) for increasing $|\theta_0|$; hence, the most stringent

requirement on the pixel size will come from θ_0 at the bounds of the integration domain. As discussed above, our analysis is restricted to $0 \leq |\theta_0|/\sigma \leq r_b$, so (10) becomes

$$\left| \frac{(r_b^2 - 1)}{24} \frac{\Delta^2}{\sigma^2} \right| \leq \epsilon, \quad (12)$$

Hence, for $\epsilon = 0.005$, the uniform grid size is $\Delta \approx 0.1\sigma$, i.e., a total of ~ 60 pixels are required per parameter-space dimension. Therefore, assuming no correlations between parameters, a grid of ≈ 3500 pixels is required to achieve 99.5% coverage of the posterior region and sub-percent net credible region identification.

In practice, the grid size can be significantly reduced in a number of ways. The grid size need not be regular; rather than requiring a fixed fractional error per pixel, we could require a fixed contribution to the absolute error, which would allow us to increase the size of pixels near the bounds of the integration volume that contain a very small fraction of the PDF but set the most stringent requirements if the fractional error criterion is used. More accurate integration can be obtained by higher-order schemes, such as Simpson's rule, lowering the minimum number of grid points necessary to achieve a given accuracy. The grid need not be rectangular, but could be an N -dimensional sphere of dimensionless radius r_b , achieving a significant volume reduction in a high-dimensional space over a cube enclosing such a sphere, as assumed above.

If the assumption of uncorrelated parameters is relaxed, (9) and (12) will still give the number of pixels required, but their relative placement needs to be altered. A misalignment between the grid and the PDF caused by correlated parameters or a non-ellipsoidal posterior PDF will reduce the validity of the previously given scaling relations. This can be overcome by either oversampling the grid; a coordinate rotation to align the grid and the PDF; or a dynamical placement of the pixels, adapting the local pixel density to a preliminary PDF estimated from a coarse grid across the parameter space. These extensions to the analysis will be investigated further in future work.

2.3. Key results

For the example BNS system described in subsection 1.1, the grid-based cumulative marginalized posterior calculation can determine the credible regions of the posterior PDF to the same accuracy as `LALInferenceMCMC` using only a small fraction ($\sim 0.1\%$) of the computational cost of the stochastic sampler. This is illustrated in Figure 1, which shows 2945 independent samples of the posterior PDF produced by `LALInferenceMCMC`, colour coded by the credible region they fall into as given by the grid-based cumulative marginalized posterior. The credible regions are compared to the fraction of MCMC samples falling into the pixels contained within them, i.e., to the credible regions as estimated by the stochastic method. The number of MCMC samples falling within a claimed credible region is governed by a binomial distribution such that the uncertainty in the fraction of samples in CR_p is $\sqrt{p(1-p)/N}$; e.g., for $N = 2945$ and $p = 0.3$, the uncertainty is 0.8% – consistent with the observed fluctuations in Figure 1. The credible




Figure 1. MCMC samples (dots) in $[M_c, \eta]$ space are colour-coded by the *cumulative marginalized posterior* credible region they fall into. The legend compares these credible regions against the fraction of MCMC samples falling into them, which corresponds to the stochastic estimate of the fraction of posterior contained within. The true parameters of the evaluated BNS system are shown at the turquoise cross (corresponding to $m_1 = 1.45M_\odot$ and $m_2 = 1.35M_\odot$).

regions, and associated uncertainties, estimated for all $p \in [0, 1]$ are shown in Figure 4 as a complement to the discrete set of credible regions displayed here.

3. Comparison with alternative methods: which approximations are warranted?

We have demonstrated that the cumulative marginalized posterior method is both accurate and computationally efficient with respect to stochastic samplers. We now explore which additional approximations can be made to simplify the analysis further; in the process, we investigate the validity of approximate techniques proposed by Baird *et al.* and Hannam *et al.* [23, 13].




Figure 2. MCMC samples (dots) in $[M_c, \eta]$ space are colour-coded by the *cumulative maximized likelihood* credible region they fall into. The legend compares these credible regions against the fraction of MCMC samples falling into them, which corresponds to the stochastic estimate of the fraction of posterior contained within. The true parameters of the evaluated BNS system are shown at the turquoise cross (corresponding to $m_1 = 1.45M_\odot$ and $m_2 = 1.35M_\odot$).

3.1. Cumulative likelihood

In the Bayesian formalism used here, it is often assumed that the majority of the information about the posterior PDF originates from the likelihood function alone, with only a weak dependence on the prior PDF. To test this assumption, we repeated the analysis performed in section 2, but without the inclusion of the prior detailed in subsection 2.1. Evaluating the cumulative marginalized likelihood across the same $[M_c, \eta]$ grid provided only negligible computational savings compared to the cumulative posterior as the overwhelming fraction of the computational cost is due to the likelihood calculations. The removal of the prior radically changed the shape of the credible regions from what was observed for the cumulative marginalized posterior in Figure 1 to the more ellipsoidal features shown in Figure 2; moreover, the credible regions computed via the cumulative marginalized likelihood method no longer match the posterior PDF as evaluated with the MCMC sampler.

We next evaluated the effect of marginalization over t_c and ϕ_c by replacing the previously used likelihood function with one which instead maximizes over t_c and ϕ_c , for both the cumulative posterior and likelihood methods. This carries greater computational savings compared to ignoring the prior, as maximizing the likelihood function reduces the computational time by $\sim 1/3$ compared to marginalizing over t_c and ϕ_c . However, replacing the correct marginalization with the computationally cheaper maximization produces credible regions which are no longer consistent with the posterior PDF estimated from MCMC calculations within the uncertainty discussed in subsection 2.3 (cf. Figure 4).

While these two simplifications, particularly the use of maximization in lieu of marginalization, can yield reductions in computational complexity, the discrepancies introduced with respect to the credible regions produced by either **LALInferenceMCMC** or the cumulative marginalized posterior, are found to be outside the required tolerance level.

3.2. Iso-Match contours and the Linear Signal Approximation

As an alternative approach for estimating credible regions and predicting parameter accuracy, Baird *et al.* [23] introduced a method, later implemented by Hannam *et al.* [13], based on the iso-match contours. This method relies on the validity of the Linear Signal Approximation (LSA) [26, 27]. Under the LSA, waveforms are assumed to vary linearly with parameters, allowing a first-order expansion

$$h(\vec{\theta}) = h_T + h_i \Delta\theta^i, \quad (13)$$

where h_i is the partial derivative of the waveform with respect to the i^{th} parameter and $\Delta\theta^i = \theta^i - \theta_T^i$. When combined with (3), this yields the likelihood function

$$L(\vec{\theta}) \propto \exp\left(-\frac{1}{2}\langle h_i | h_j \rangle \Delta\theta^i \Delta\theta^j\right), \quad (14)$$

assuming $n \equiv \vec{0}$, expressed as a multivariate Gaussian centred at the true parameters $\vec{\theta}_T$ with covariance matrix $\langle h_i | h_j \rangle^{-1}$.

The method of Baird *et al.* uses the waveform match M between waveforms h_T and $h(\vec{\theta})$ defined as

$$M = \max_{t_c, \phi_c} \frac{\langle h_T | h(\vec{\theta}) \rangle}{\sqrt{\langle h_T | h_T \rangle \langle h(\vec{\theta}) | h(\vec{\theta}) \rangle}}, \quad (15)$$

again maximizing over t_c and ϕ_c , as a proxy for the likelihood function. By assuming that the LSA is valid, Baird *et al.* approximated credible region boundaries as contours of constant match,

$$M_p = 1 - \frac{\chi_N^2(1-p)}{2\rho^2}, \quad (16)$$

via an N -dimensional χ^2 distribution where N is again the number of dimensions of the parameter space remaining after maximization. In addition to using maximization in




Figure 3. MCMC samples (dots) in $[\mathcal{M}_c, \eta]$ space are colour-coded by the credible region they fall into as determined by the *iso-match contour* bounding them. The legend compares these credible regions against the fraction of MCMC samples falling into them, which corresponds to the stochastic estimate of the fraction of posterior contained within. The true parameters of the evaluated BNS system are shown at the turquoise cross (corresponding to $m_1 = 1.45M_\odot$ and $m_2 = 1.35M_\odot$).

lieu of marginalization, the validity of this approximation relies on the Gaussianity of the posterior, and does not include a priori information.

Calculating the matches given for each pixel in the same $[\mathcal{M}_c, \eta]$ grid as used in Figure 1 and 2, we defined credible regions as the pixels bounded by the iso-match contour in (16). Figure 3 compares credible regions given by iso-match contours against estimates from the fraction of MCMC samples falling within those contours; the differences between the two are statistically significant.

While the matches used by this method are calculated exactly, the restrictions implied by the LSA will lead the method to fail if the posterior PDF under investigation exhibits even a moderate level of non-Gaussianity. This can originate in the likelihood itself, or from the neglected contribution of the prior. This becomes clear for the BNS system evaluated here from the high degree of similarity between Figures 2 and 3, indicating the validity of the LSA for this system. In this particular case, the Gaussianity of the likelihood in $[\mathcal{M}_c, \eta]$ space means that the posterior would have been Gaussian

if the priors were flat in $[\mathcal{M}_c, \eta]$ space, so the method could have performed relatively well; it does not perform well for flat priors in $[m_1, m_2]$ space, as indicated by Figure 4 (see below), because the posterior in this case is no longer Gaussian.

3.3. Comparasion

We compare all of the grid-based methods described above in Figure 4. We show the differences between the fraction of MCMC samples contained within the various methods for estimating credible regions corresponding to credible level p , and the value of p , for $p \in [0, 1]$. Perfect agreement would correspond to a horizontal line at a deviation of zero. However, the finite number of stochastic MCMC samples from `LALInferenceMCMC` leads to statistical fluctuations in the deviation; their expected magnitude is indicated by an error ellipse (see subsection 2.3) corresponding to one- σ fluctuations.

The cumulative posterior method, using a likelihood function marginalized over t_c and ϕ_c , successfully estimates credible regions (apparent deviations at $p < 0.3$ could be statistical, or may be due to the need for sub-pixel resolution to resolve such small credible regions). The comparison also further solidifies both the validity of the LSA for this system and the effect of the prior on the ability to recover consistent credible regions with respect to `LALInferenceMCMC`.

4. Conclusions and future directions

We have evaluated several grid-based methods for approximating the parameter credible regions for a CBC event, and compared these to regions estimated by the stochastic sampler `LALInferenceMCMC`. We found that evaluating the cumulative posterior on a relatively low-density grid allowed us to estimate credible regions to within the statistical uncertainty of the stochastic sampler at a small fraction of the computational cost ($\sim 0.1\%$), while marginalizing over the time and phase parameters and incorporating an arbitrary prior.

On the other hand, ignoring the prior or maximizing over t_c and ϕ_c instead of marginalizing over them introduces a discrepancy in the recovered credible regions with respect to `LALInferenceMCMC`. In addition, we have demonstrated that the iso-match method proposed by Baird *et al.* is overestimating the credible regions in $[\mathcal{M}_c, \eta]$ space compared to a full Bayesian analysis.

The analysis has been performed on a binary observed in one detector at a fixed overhead and optimally oriented position and at a fixed distance giving $\rho = 12$. Compared to an analysis of the same binary using a three-detector observation comprising data from the two LIGO observatories and the Virgo observatory (all operating at an the same sensitivity as assumed in subsection 1.1), where extrinsic parameters describing sky location, inclination, orientation and distance are allowed to vary, the recovered two-dimensional credible region in $[\mathcal{M}_c, \eta]$ space, is not significantly altered (see Figure 5). The analysis is implemented with the same $[\mathcal{M}_c, \eta]$ grid




Figure 4. Difference between the fraction of MCMC samples from the posterior PDF falling into credible regions predicted by the various methods described above and the expected credible level p , as a function of p . Continuous relations for $p \in [0, 1]$ of the data in in Figure 1, 2 and 3 are shown by five curves, corresponding, from the top down, to the cumulative marginalized likelihood, cumulative maximized likelihood, iso-match contours, cumulative marginalized posterior, and cumulative maximized posterior. The filled ellipse indicates the expected 68% level of fluctuation in the fraction of MCMC samples due to the finite number of samples.

as in previous figures, a grid designed for a two-dimensional analysis as described in subsection 2.2. The discrepancy, if any, introduced by allowing for eventual correlation caused by the inclusion of extrinsic parameters is found to be within the statistical uncertainty from the limited number of samples from `LALInferenceMCMC` (c.f. subsection 2.3).

The cumulative marginalized posterior method can easily be extended to include other parameters, especially spin [13], but as the computational cost scales exponentially with the number of parameters, the cost quickly approaches and exceeds the computational requirements of the stochastic sampler implemented in `LALInference`. Our simple grid-based sampling implementation would be computationally competitive with the stochastic sampling for parameter spaces with up to *four* non-marginalized dimensions. However, the computational cost of the grid-based sampler could be further




Figure 5. MCMC samples (dots), now drawn from a 9-dimensional analysis, in $[\mathcal{M}_c, \eta]$ space are colour-coded by the *cumulative marginalized posterior* credible region they fall into, using the same two-dimensional grid as in previous figures. The legend compares these credible regions against the fraction of MCMC samples falling into them, which corresponds to the stochastic estimate of the fraction of posterior contained within, now allowing for effects of correlation against extrinsic parameters. The true parameters of the evaluated BNS system are shown at the turquoise cross (corresponding to $m_1 = 1.45M_\odot$ and $m_2 = 1.35M_\odot$).

reduced through more efficient grid placement and more accurate integration algorithms. Moreover, while standard stochastic samplers such as `LALInferenceMCMC` are serial processes, all pixels in the sampling grid are completely independent, therefore trivially allowing for massive parallelization of the cumulative marginalized posterior analysis.

Through these properties we envision the grid-based sampling method, using a cumulative marginalized posterior, to be implemented as a low-latency parameter estimation tool for the intrinsic parameters of a CBC candidate event, similar to the implementation of `BAYESTAR` for the extrinsic parameters [10, 28]. In practice, we won't know the true signal parameters which are needed for efficient and accurate grid placement. However, we can take the parameters of the highest-match template from the search pipelines used for the detection of CBC events [29, 30], which by design of the template banks will generally have $M > 0.97$ [31], as the central point of the grid. The

size of the grid can be initially estimated by comparing the SNRs reported by adjacent templates in the template bank. Subsequently, the grid can be adaptively refined: while we used a grid with uniformly spaced pixels for this study, the computational cost could be reduced further by implementing a non-uniform grid. The reduction in the absolute number of pixels required within the grid can enable the inclusion of additional non-marginalized dimensions while retaining the computational competitiveness against the stochastic sampling methods.

Moreover, predictive methods such as the iso-match method [23, 13] or the effective Fisher matrix approach [32, 33] can be used in combination with the grid-based cumulative marginalized posterior technique to provide PDF estimates for dynamically laying out the grid. The cumulative marginalized posterior could also be implemented as a jump proposal for the intrinsic parameters as part of the stochastic samplers in **LALInference**. These possibilities will be explored further in future work.

In addition, it is important to note that even though the method of parameter estimation with the grid-based sampling using a cumulative marginalized posterior has been presented here in the context of gravitational-wave astrophysics, the method itself is completely general and can be effective whenever the dimensionality of the parameter space is sufficiently small to make it competitive with stochastic samplers.

Acknowledgments

We are grateful for computational resources provided by the Leonard E Parker Center for Gravitation, Cosmology and Astrophysics at University of Wisconsin-Milwaukee. We received useful comments and suggestions from a number of colleagues at the University of Birmingham, including Christopher Berry, Walter Del Pozzo, Alberto Vecchio, and John Veitch. We are grateful to Richard O’Shaughnessy and Chris Pankow for careful reading of the manuscript. This work was partly supported by the Science and Technology Facilities Council and the Leverhulme Trust. IM acknowledges the hospitality of the Monash Center for Astrophysics, supported by a Monash Research Acceleration Grant (PI Y. Levin). This is LIGO document number P1400251.

References

- [1] G. M. Harry and the LIGO Scientific Collaboration. Advanced LIGO: the next generation of gravitational wave detectors. *Class. Quantum Gravity*, 27(8):084006, April 2010, arXiv:1103.2728.
- [2] The LIGO Scientific Collaboration. Advanced LIGO. *ArXiv e-prints*, November 2014, 1411.4547.
- [3] Virgo Collaboration. Advanced virgo baseline design. Virgo Technical Report VIR-0027A-09, 2009. <https://tds.ego-gw.it/itf/tds/file.php?callFile=VIR-0027A-09.pdf>.
- [4] F Acernese, M Agathos, K Agatsuma, et al. Advanced virgo: a second-generation interferometric gravitational wave detector. *Classical and Quantum Gravity*, 32(2):024001, 2015.
- [5] LIGO Scientific Collaboration, Virgo Collaboration, J. Aasi, et al. Prospects for Localization

- of Gravitational Wave Transients by the Advanced LIGO and Advanced Virgo Observatories. ArXiv e-prints, April 2013, 1304.0670.
- [6] J. M. Weisberg and J. H. Taylor. The Relativistic Binary Pulsar B1913+16: Thirty Years of Observations and Analysis. In F. A. Rasio and I. H. Stairs, editors, Binary Radio Pulsars, volume 328 of Astronomical Society of the Pacific Conference Series, page 25, July 2005, astro-ph/0407149.
 - [7] M. Kramer, I. H. Stairs, R. N. Manchester, et al. Tests of general relativity from timing the double pulsar. Science, 314(5796):97–102, 2006, <http://www.sciencemag.org/content/314/5796/97.full.pdf>.
 - [8] J. Abadie, B. P. Abbott, R. Abbott, et al. TOPICAL REVIEW: Predictions for the rates of compact binary coalescences observable by ground-based gravitational-wave detectors. Classical and Quantum Gravity, 27(17):173001, September 2010, 1003.2480.
 - [9] D. A. Brown, I. Harry, A. Lundgren, and A. H. Nitz. Detecting binary neutron star systems with spin in advanced gravitational-wave detectors. Phys. Rev. D, 86(8):084017, October 2012, 1207.6406.
 - [10] Leo P. Singer, Larry R. Price, Ben Farr, et al. The first two years of electromagnetic follow-up with advanced ligo and virgo. The Astrophysical Journal, 795(2):105, 2014.
 - [11] C. P. L. Berry, I. Mandel, H. Middleton, et al. Parameter estimation for binary neutron-star coalescences with realistic noise during the Advanced LIGO era. ArXiv e-prints, November 2014, 1411.6934.
 - [12] M. Favata. Systematic Parameter Errors in Inspiring Neutron Star Binaries. Physical Review Letters, 112(10):101101, March 2014, 1310.8288.
 - [13] Mark Hannam, Duncan A. Brown, Stephen Fairhurst, Chris L. Fryer, and Ian W. Harry. When can gravitational-wave observations distinguish between black holes and neutron stars? The Astrophysical Journal Letters, 766(1):L14, 2013.
 - [14] Alessandra Buonanno, Bala R. Iyer, Evan Ochsner, Yi Pan, and B. S. Sathyaprakash. Comparison of post-newtonian templates for compact binary inspiral signals in gravitational-wave detectors. Phys. Rev. D, 80:084043, Oct 2009.
 - [15] Advanced LIGO anticipated sensitivity curves. <https://dcc.ligo.org/LIGO-T0900288/public>.
 - [16] B. P. Abbott, R. Abbott, R. Adhikari, et al. Search for gravitational waves from low mass binary coalescences in the first year of LIGO’s S5 data. Phys. Rev. D, 79(12):122001, June 2009, 0901.0302.
 - [17] LSC Algorithm Library software packages LAL, LALWRAPPER, and LALAPPS.
 - [18] J. Veitch, V. Raymond, B. Farr, et al. Parameter estimation for compact binaries with ground-based gravitational-wave observations using the lalinference software library. Phys. Rev. D, 91:042003, Feb 2015.
 - [19] J. Aasi, J. Abadie, B. P. Abbott, et al. Parameter estimation for compact binary coalescence signals with the first generation gravitational-wave detector network. Phys. Rev. D, 88:062001, Sep 2013.
 - [20] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of State Calculations by Fast Computing Machines. J. Chem. Phys., 21:1087–1092, June 1953.
 - [21] W. K. Hastings. Monte carlo sampling methods using markov chains and their applications. Biometrika, 57(1):97–109, 1970.
 - [22] T. Sidery, B. Aylott, N. Christensen, et al. Reconstructing the sky location of gravitational-wave detected compact binary systems: Methodology for testing and comparison. Phys. Rev. D, 89(8):084060, April 2014, 1312.6013.
 - [23] Emily Baird, Stephen Fairhurst, Mark Hannam, and Patricia Murphy. Degeneracy between mass and spin in black-hole-binary waveforms. Phys. Rev. D, 87:024035, Jan 2013.
 - [24] Will M. Farr, Ilya Mandel, and John Veitch. Speeding up parameter estimation on gravitational waves from coalescing compact binaries: marginalizing over time and phase, 2015. In preparation.

- [25] P. C. Mahalanobis. On the generalised distance in statistics. *Proceedings of the National Institute of Sciences of India*, 2(1):49–55, April 1936.
- [26] Michele Vallisneri. Use and abuse of the fisher information matrix in the assessment of gravitational-wave parameter-estimation prospects. *Phys. Rev. D*, 77:042001, Feb 2008.
- [27] C. L. Rodriguez, B. Farr, W. M. Farr, and I. Mandel. Inadequacies of the Fisher information matrix in gravitational-wave parameter estimation. *Phys. Rev. D*, 88(8):084013, October 2013, 1308.1397.
- [28] L. P. Singer. *The needle in the hundred square degree haystack: The hunt for binary neutron star mergers with LIGO and Palomar Transient Factory*. PhD thesis, California Institute of Technology, January 2015. <http://resolver.caltech.edu/CaltechTHESIS:12102014-223122387>.
- [29] S. Babak, R. Biswas, P. R. Brady, et al. Searching for gravitational waves from binary coalescence. *Phys. Rev. D*, 87(2):024033, January 2013, 1208.3491.
- [30] Kipp Cannon et al. Toward early-warning detection of gravitational waves from compact binary coalescence. *The Astrophysical Journal*, 748(2):136, 2012.
- [31] B. J. Owen and B. S. Sathyaprakash. Matched filtering of gravitational waves from inspiraling compact binaries: Computational cost and template placement. *Phys. Rev. D*, 60(2):022002, July 1999, gr-qc/9808076.
- [32] Hee-Suk Cho, Evan Ochsner, Richard O’Shaughnessy, Chunglee Kim, and Chang-Hwan Lee. Gravitational waves from black hole-neutron star binaries: Effective fisher matrices and parameter estimation using higher harmonics. *Phys. Rev. D*, 87:024004, Jan 2013.
- [33] Richard O’Shaughnessy, Ben Farr, Evan Ochsner, et al. Parameter estimation of gravitational waves from nonprecessing black hole-neutron star inspirals with higher harmonics: Comparing markov-chain monte carlo posteriors to an effective fisher matrix. *Phys. Rev. D*, 89:064048, Mar 2014.